

Anticipating the Effects of Technological Change: A New Era of Dynamics for Human Factors

David Woods

Institute for Ergonomics
The Ohio State University, USA

Sidney Dekker

Linköping Institute of Technology, Sweden

Abstract

Human factors studies the intersection between people, technology and work, with the major aim to find areas where design and working conditions produce human error. It relies on the knowledge base and research results of multiple fields of inquiry (ranging from computer science to anthropology) to do so. Technological change at this intersection (1) redefines the relationship between various players (both humans and machines), (2) transforms practice and shifts sources of error and excellence, and (3) often drives up operational requirements and pressures on operators. Human factors needs to predict these reverberations of technological change before a mature system has been built in order to steer design into the direction of cooperative human-machine architectures. The quickening tempo of technology change and the expansion of technological possibilities has largely converted our traditional shortcuts for access to a design process (task analysis, guidelines, verification & validation studies, etc.) into oversimplification fallacies that retard understanding, innovation, and ultimately, human factors' credibility. There is an enormous need for the development of techniques that gain empirical access to the future--that generate human performance data about systems which have yet to be built.

Introduction

The Reverberations of Technology Change on Fields of Practice

Human Factors as a field is based on observing people at work. To the degree that we abstract patterns from this process of observation, one can view Human Factors as the body of work that describes how technology and organizational change transforms work in systems (Woods, Christoffersen and Tinapple, 2000). When we observe the introduction of new technology and systems into a field of practice, we see how the change represents new ways of doing things; i.e., it does not preserve the old ways with the simple substitution of

one medium for another (e.g., paper for computer-based). The reality of technology change is **transformation** and **adaptation** (Carroll's task-artifact cycle; Carroll and Rosson, 1992). The idea that new technology can be introduced as a simple substitution of machines for people—preserving the system though improving the results—is a persistent oversimplification fallacy: the substitution myth. In actuality, adding or expanding the machine's role changes the cooperative architecture and changes the human's role (Sarter, Woods and Billings, 1997).

Technology change is an intervention into an ongoing field of activity (Winograd and Flores, 1987; Flores, Graves, Hartfield, and Winograd, 1988). The studies of this process point out:

“It [new technology] alters what is already going on—the everyday practices and concerns of a community of people—and leads to a resettling into new practices.” (Flores et al., 1988, p. 154).

“New tools alter the tasks for which they were designed, indeed alter the situations in which the tasks occur and even the conditions that cause people to want to engage in the tasks.” (Carroll and Campbell, 1988, p. 4).

The review of the impact of new technology on the Desert Storm operation summarizes the general pattern remarkably well (Cordesman and Wagner, 1996, p.25):

Much of the equipment deployed ... was designed to ease the burden on the operator, reduce fatigue, and simplify the tasks involved in operations. Instead, these advances were used to demand more from the operator.

Almost without exception, technology did not meet the goal of unencumbering the personnel operating the equipment ... systems often required exceptional human expertise, commitment, and endurance.

... there is a natural synergy between tactics, technology, and human factors ... effective leaders will exploit every new advance to the limit. As a result, virtually every advance in ergonomics was exploited to ask personnel to do more, do it faster and do it in more complex ways.

... one very real lesson is that new tactics and technology simply result in altering the pattern of human stress to achieve a new intensity and tempo of operations.
[edited to rephrase military referents generically]

This statement could have come from studies of the impact of technological and organizational change in health care or air traffic management or many other areas undergoing change today (e.g., Cook and Woods, 1996a; 1996b; Obradovich and Woods, 1996; Smith et al., 1998; Dekker and Woods, 1999). The pattern is that technology change transforms operational and cognitive systems:

- new roles emerge,
- what is canonical (routine) and what is exceptional changes,
- the kinds of erroneous actions and assessments that can be expected change,
- the paths to failure change.

People in their various roles adapt to achieve goals and avoid failure:

- by tailoring in the face of poor human-computer cooperation,
- by developing and modifying failure sensitive strategies.

Yet performance pressures on the overall system (greater efficiency or throughput) tend to push practitioners back to the edge of the performance envelope rather than taking the benefits of the changes in increased safety margin or lower workload. As a result, surprises occur in the form of accidents (fundamentally surprising new paths to failure) and in the form of negative side effects of the change unanticipated by designers—“automation surprises” (Sarter and Amalberti, 2000 and Dekker and Hollnagel, 1999 for the case of cockpit automation).

Overall, the studies show that when “black box” new technology (and accompanying organizational change) hits an ongoing field of practice the pattern of reverberation is (Woods et al., 1994, chapter 5):

- New capabilities, which increase demands and create new complexities such as increased coupling across parts of the system and higher tempo of operations.
- Additional new complexities when technological possibilities are used clumsily,
- Adaptations by practitioners because they are responsible to meet operational goals,
- The complexities and adaptations are surprising, unintended side effects of the design intent,
- Failures occasionally break through these adaptations because the adaptations are poor or brittle and because other circumstances arise which help move conditions toward failure,
- The adaptations by practitioners hide the complexities from designers and reviewers after-the-fact who judge failures to be due to human error.

The results illustrate a more general law of adaptive systems that has been noted by many researchers (e.g., Rasmussen, 1986; Hirschhorn, 1997)—*the law of stretched systems*:

every system is stretched to operate at its capacity; as soon as there is some improvement, for example in the form of new technology, it will be exploited to achieve a new intensity and tempo of activity.

Under pressure from performance and efficiency demands, advances will be used to ask operational personnel to do more, do it faster or do it in more complex ways. To stem this pressure and guide technological or other interventions in more fruitful directions, one of the founding slogans of Cognitive Systems Engineering has been—*adaptations directed at*

coping with complexity (Rasmussen and Lind, 1981; Woods, 1988; Woods, Christoffersen and Tinapple, 2000).

A New Era of Dynamics Challenges Traditional Shortcuts

In the last era, a slower pace of change allowed Human Factors to adopt basic simplification strategies—a propensity to simplify by converting a dynamic process into a static snapshot and a propensity to simplify by converting multiple-factor, interconnected processes into assessing the state of few independent things (Woods and Tinapple, 1999). However, the quickening tempo of technology change and the expansion of technological possibilities has converted these shortcuts into oversimplification fallacies that retard understanding, innovation, and ultimately, human factors' credibility.

Given the dynamics of system change in the era of rapidly changing technological possibilities, success will come to those who can *predict* the transformations, changing roles, etc., predict the kinds of adaptations practitioners develop to cope with new complexities, and predict the situations which will challenge these strategies to anticipate where errors and failure may emerge in the future (Rasmussen, 2000; Corker, 2000). Second, success will come to those who have the ability to use these predictions early in the design process to avoid the negative unintended side effects of technology change. The pace of change created by expanding technological possibilities demands a dynamics of people, technology and work. Observing these dynamics, modeling these dynamics, and learning to gently steer the processes of change, given the stakes for the multiple parties affected by these changes, is the challenge of the new era.

Developing a dynamics of people, technology and work raises several challenges. The phenomena of interest lie at the intersection of traditional fields of inquiry overlapping and connecting technological and behavioral sciences, individual and social perspectives, the laboratory and the field, design activity and empirical investigation, theory and application. As a result, successful inquiry will depend on a new class of “inter-cultural brokers.” Such a class of Human Factors professionals would recognize that a new design is not simply an object, but also a hypothesis about how technological change will transform practice and shape how people will adapt (Woods, 1998). This means that designers also function as experimenters--the adaptive response of people and organizations to new systems tests the hypotheses about what would be useful embodied by those particular prototypes or products. Balancing the parallel status of new designs as objects to be realized and as hypotheses about how change affects cognitive and collaborative work requires new processes of inquiry and of development (Woods, Christoffersen and Tinapple, 2000; Sanders, 2000). This also means that to design effectively requires an understanding of the field of practice the design is intended to support. In a simplification shortcut, we assumed in the past that studying the current world would meet that requirement despite the fact that the introduction of the new systems would change what it meant to practice in that setting.

The tempo and ramification of change is outpacing conventional forms of inquiry we have relied on previously in Human Factors and in related areas. As in all dynamic processes, there is hysteresis, as institutions, conventional practices, and conventional beliefs linger on despite being ill-suited, inappropriate, or erroneous with respect to the tasks at hand. The mismatch between old habits of mind and method and exciting new opportunities results in strong tensions as the old guard “forts up” to vigorously defend irrelevant turf, wielding tools and relics of a bygone era.

From Evaluating New Systems to Generating New Concepts

In the past, Human Factors worked with many industries to assess the impact of new systems intended to aid human performance. Stakeholders frequently saw our role as providing a simple up or down result—does this particular system or technology help significantly or not?—in verification and validation evaluations (V&V) or as tuning the system as a late step prior to release (usability testing). Inevitably, Human Factors, when institutionalized in these ways, has found itself pushed to the tail of the design process (Woods and Tinapple, 1999).

We find over and over again that providing such empirical testing roles provide too little information, too late in the design process, at too great a cost. There are multiple degrees of freedom in using new technology to design systems, but late testing studies are not able to tell developers how to use those degrees of freedom to create useful and desirable systems. The problem in design today is not can we build it, but rather what would be useful to build given the wide array of possibilities new technology provides.

New systems and technology are not unidimensional, but multi-faceted, so problems of credit assignment easily become overwhelming in late testing studies. Introducing new technology is not manipulating a single variable, but a change that reverberates throughout a system transforming judgments, roles, relationships, and weightings on different goals.

This process, the task-artifact cycle, creates the *envisioned world problem* for research and design (Dekker and Woods, 1999; Hoffman and Woods, 2000): how can results of studies and analyses that characterize cognitive and cooperative activities in the current field of practice inform or apply to the design process, *since the introduction of new technology will transform the nature of practice*, what it means to be an expert, and the paths to failure?

A variety of factors push testing studies too late in the design process, especially given their great cost, to provide useful input. By the time that the results are available, the design process has committed to certain design concepts and implementation directions. These sunk costs make it extremely difficult to act on what is learned from evaluation studies late in the process.

The standard retort to such difficulties is to point to technology for rapid prototyping of possible designs. And rapid prototyping technology is indeed a critical prerequisite, but it is not sufficient. Alone and given the resource pressures and limited time horizon of all real

development projects, rapid prototyping technology ends up only speeding up the same old process with the same difficulties—as we have quipped before, “with rapid prototyping (where prototypes function only as partially refined final products), we make the same mistakes, only faster” (Woods et al., 1996).

To deal with the task-artifact cycle and the envisioned world problem, we need something more to help generate what would be useful in a dynamic, participatory process of change, pressure, and adaptation. This need has led to a complete shift in emphasis in resource investment in design processes, from relying on testing near the end of development iterations and cycles to early, generative techniques such as ethnography and participatory design to help envision new possibilities and directions (Greenbaum and Kyng, 1991; Carroll and Rosson, 1992; Smith et al., 1998; Woods et al., 2000; Sanders, 2000).

The Envisioned World Problem

The scope and pace of change severely limits the usefulness of our traditional tactics. The dynamics of people, technology and work demands we face up to the envisioned world problem. Since the introduction of new technology transforms the nature of practice, we must innovate techniques to answer questions like:

- How does one envision or predict the relation of technology, cognition and collaboration in a domain that doesn't yet exist or is in a process of becoming?
- How will envisioned technological change shape cognition and collaboration?
- How will practitioners adapt artifacts, given mismatches to the actual demands and pressures they experience, to meet their own goals?
- How can we predict the changing nature of expertise and new forms of failure as the workplace changes?
- How will design processes create new tools that useful and robust since there are limits to predictions of a co-evolutionary process?

All parties to a field of practice and an episode of change envision possible futures. The investment of energy and resources in a development project is justified in part on the basis of its presumed benefits for human performance. This means prototypes and products embody a *hypothesis* that changing parts of an operational world will carry benefits for human cognitive and collaborative activities. But the actual, rather than presumed, impact of new technology is usually quite surprising, unintended, and even counterproductive. The surprises occur because of the co-evolutionary process at work: in the face of new complexities and capabilities, human users actively adapt the technology provided to them, tailoring it and their strategies to the immediate tasks at hand in a locally pragmatic way.

Thus, one critical issue is how to ground that envisioning to lawful factors at the intersection of people, technology and work. As a result, designers need to adopt the attitude of an experimenter trying to understand and model the interactions of task demands, artifacts, cognition, collaboration across agents, and organizational context (Woods, 1998). An experimental stance means that designers need to

- recognize that design concepts represent hypotheses or beliefs about the relationship between technology and cognition/collaboration,
- subject these beliefs to empirical jeopardy by a search for disconfirming and confirming evidence,
- recognize that these beliefs about what would be useful are tentative and open to revision as we learn more about the mutual shaping that goes on between artifacts and actors in a field of practice.

This kind of experimental stance for design is needed to make a difference in developing useful systems and it represents a challenge to traditional modes of experimental research on human behavior at work.

Predicting Post-Conditions of Technology Change

If Human Factors is grounded on observing and describing how technology and organizational change transforms work in systems, then the test of success in the future will be our ability to

- (a) anticipate unintended effects—how design can err in the sense that design will create conditions which will produce certain kinds of undesirable behaviors by people operating in that system.
- (b) innovate—using the information about dynamics of change and adaptation to create new and reusable design possibilities.

There are four components to a dynamics of Human Factors needed to meet these criteria:

- **data:** observe the processes of change and adaptation to abstract patterns.
- **models:** generalizations about the essential variables that drive this dynamic process which explain the observed patterns across specific cases.
- **prediction:** given proposed new technology and the goals and expectations of advocates for that change, what are the likely changes, adaptations, vulnerabilities, and other reverberations that will play out.
- **design:** using the predictions to redirect design into more fruitful, reusable directions.

The first step, to cope with the dynamic, co-evolutionary process of technology change and adaptation through use in context, is the need for data in the form of observations of those processes. Human Factors could engage more in observing episodes of change and abstracting patterns about how organizational and technology change transforms cognitive and collaborative demands and activities, and how in turn people adapt to those changes. These observations and patterns can then drive explanation building and modeling—generalizations about the dynamics of change and adaptation.

The third step is to use the research base of data and models on transformation and adaptation to anticipate consequences of technological interventions in specific settings. These predictions, grounded on past results, drive investigations of how possible or proposed technology changes will transform roles, demands and activities, and where it will insert new vulnerabilities into the operational world. The goal of this process is to identify

post-conditions of technology change. Look ahead is always tentative and tenuous, but the best tool to support look ahead is a good model that captures the essential variables in the dynamics of change and adaptation.

Finally, the ultimate purpose is to stimulate design and innovation—sparking inventiveness to discover new ways to use technological possibilities to enhance human performance, to identify leverage points, and to minimize unanticipated side effects. The possibilities of technology afford designers great degrees of freedom. The possibilities seem less constrained by questions of feasibility and more by concepts about how to use the possibilities skillfully to meet operational and other goals. The pressing need is to determine what will be useful as we create future possibilities, remembering that we are participants with other stakeholders and problem holders in the processes of change in that field of practice.

Constraints on Envisioned Operations

The relationship between technology and cognition and collaboration is challenging to trace out because of the dynamic interplay at work—practitioners adapt to difficulties, re-shaping artifacts to function as tools to meet the demands of the field of activity. The dynamic of this process creates the moving target for development that we call the Envisioned World Problem (Dekker, 1996; Dekker and Woods, 1999).

Changing demands, pressures and resources lead stakeholders to envision and advocate for new possibilities. Envisioning future operations is ubiquitous in work on advancing the baseline of technology and in advanced development projects in specific areas. However, envisioned operation concepts have two basic properties:

- *plurality* – there are multiple versions of how the proposed changes will effect the character of the field of practice in the future; and
- *underspecification* – each envision concept is vague on many aspects of what it would mean to function in that field of practice in the future; in other words, each is a simplification, or partial representation of what it will mean to practice when that envisioned world becomes concrete.

In envisioning future operations, different groups of stakeholders (regulators, technology vendors, practitioners) have different conceptions or visions of what the future world will be like. Since each comes from only the single angle of one group of practitioners or stakeholders, these views are necessarily partial representations, or simplifications of reality. Cumulatively, there is a loosely coupled collection of visions for the future; visions that vary greatly in perspective and degree of detail.

In addition, they may be driven in part by parochialism or advocacy, as different groups of practitioners will have different stakes in what the future world will - or must - have in store for them, their positions, influence, job security, status,

or roles. The fact that nobody is bound by tangible technologies or procedures can amplify the gaps and simplifications.

The question then is—can design anticipate the full range of potential effects, or *post-conditions*, of the change? Usually, technology change produces unintended and sometimes negative side effects in addition to new capabilities. Thus we are part of a dynamic process which we wish to understand and influence – a dynamic process of technology change generating a new set of capabilities and complexities, leading to adaptations by stakeholders, producing a changing mix of success and failure.

Since envisioned modes of operation are a prediction about the effects of change on people, technology and work, they can have two other properties:

- *ungrounded* – envisioned concepts can easily be disconnected or even contradict from the research base on the actual consequences of the changes on people, technology and work.
- *overconfident* – advocates are miscalibrated and overconfident that if the systems envisioned can be realized the predicted consequences and only the predicted consequence will occur.

The envisioned world problem demands that we develop means to ground predictions on relevant empirical results abstracted from observations in context. Understanding the dynamic process of change and adaptation will lead to better control of the process – essentially an innovation process at the intersection of people, technology and work. Armed with knowledge about the dynamics of change and adaptation, we can address potential side effects at a time when intervention is less difficult and less expensive (because the field of practice is already in a period of change and systems development is in the process of creating tangible objects).

Assessing Post-Conditions of Change

There are several ways to address the envisioned world problem as a constraint on post-condition analysis. Studies of the current world can contribute, in part, to the degree that it models the demands of the field of practice—what factors make problems hard, what are complicating factors that push situations beyond textbook cases. Functional models of how the underlying process works can identify demands that any set of cognitive agents or strategies must accommodate (Roth and Mumaw, 1995). Artifact-based methods, where prototypes function as a kind of experimental probe and tool for discovery when placed in the hands of practitioners in meaningful scenarios, is another way to deal with the envisioned world problem.

Another approach is developing scenarios that capture fundamental demands of the field of activity can be used to support walkthroughs and simulations of possible future operational concepts (Carroll, 2000). The future incident technique is based on this (Dekker, 1996). This technique aims to develop a failure or near miss that could happen given a general view

of how the envisioned world will work. The future incident is based on two things: first on our knowledge of technology-independent vulnerabilities or challenges in the field of practice (e.g. . in the case of future air traffic management, merging or crossing streams of traffic, depressurizations, clear air turbulence etc). Second, it is based on knowledge of classic design errors—cases where new technology shaped cognition and collaboration in ways that produced problems such as clumsy automation, coordination surprises or other difficulties that contributed to incidents or accidents (Norman, 1988). The scenario designers look at different views of the envisioned world, searching for places where design errors would negatively influence cognition or collaboration of future practitioners. In one variation, the resulting future incident is packaged in the format of actual incident reporting formalisms used in that field of practice: different kinds of participants in the current system then view these full incident reports about their possible future world. In another, participants are given the initial conditions of a particular future situation, after which a system perturbation is thrown in to create a sample of future problem solving situations (see Dekker and Woods, 1999).

In both cases, the use of concrete scenarios to anchor participants in the details of coordination, communication, decision making and knowledge exchange necessary to handle the situation successfully is critical to this method. The incidents are illustrations of where a future architecture may be vulnerable or how it may break down, thus inviting practitioners and developers alike to think critically about the requirements for effective problem-solving in the envisioned world. Note how this inverts human factors involvement in system development: no longer is the emphasis on verifying for developers that a particular approach may work. Instead the focus is on exploring cognitive pressure points; those areas where situational demands or features will outwit the problem-solving resources envisioned to be in the hands of future practitioners (Weick, 1979; DeGeus, 1988).

Techniques to tackle the envisioned world problem shift the dominant sources of validity in research on human performance as compared to more developed or existing operational environments. In the latter, face validity of a simulation tool or experimental set-up is often thought to provide much of the requisite mapping between test situation and target world. In contrast, in research on envisioned worlds, validity (or perhaps better, authenticity; Woods et al., 2000) derives from (1) the extent to which problems-to-be-solved in the test situation represent the vulnerabilities and challenges that exist in the target world, and (2) the extent to which real problem-solving expertise is brought to bear by the study participants (Klein et al., 1993). Studies into human performance in envisioned worlds can rate high on both of these measures by (a) creating future incidents, and (b) involving real practitioners who have been prepared for their future roles. In other words, these studies must investigate real practitioners caught up in solving real domain problems. The future incident technique, as one solution to the envisioned world problem, is needed because Human Factors must identify technological complexities and their consequences early in the design process. In this way, system development can be steered towards more fruitful, cooperative channels, breaking all participants in the development process out of their usual frames of reference.

Conclusion

The pace of technological and organizational change is rapid. Consequences of change are profound in hindsight, but difficult to anticipate. Past simplifications in how Human Factors examined people, technology and work are ineffective in the face of the pressure to help understand and direct such powerful forces. Creating the new era of dynamics for Human Factors is a paradigm shift:

- that remakes what, where, and how we observe people at work,
- that demands new kinds of explanations for the phenomena of interest,
- that challenges us to predict the reverberations of envisioned changes, and
- that invites us to be part of the innovation process that creates future possibilities.

The foundation for the future is observations of how technological and organizational changes transform cognitive and collaborative activities and demands, and how people in turn adapt to those changes.

Acknowledgments

Work for this article was supported in part by grant 93-G-036 from the U.S. Federal Aviation Administration and a grant from the Flight Safety Department at the Swedish Civil Aviation Authority.

References

- Carroll, J.M. and Campbell, R.L. (1988). Artifacts as psychological theories: The case of human-computer interaction. IBM Research Report RC 13454, Watson Research Center, Yorktown Heights, NY.
- Carroll, J.M. and Rosson, M. B. (1992). Getting around the task-artifact cycle: How to make claims and design by scenario. *ACM Transactions on Information Systems*. 10, 181-212.
- Carroll, J.M. (2000). Making Use: Scenario-Based Design of Human-Computer Interactions. MIT Press, Cambridge, MA.
- Cook, R.I., and Woods, D.D. (1996a). Adapting to new technology in the operating room. *Human Factors*, 38(4), 593-613.
- Cook R.I. and Woods. D.D. (1996b). Implications of automation surprises in aviation for the future of total intravenous anesthesia (TIVA). *Journal of Clinical Anesthesia*, 8:29s-37s.

- Corker, K. M. (2000). Cognitive Models and Control: Human and System Dynamics in Advanced Airspace Operations. In N. Sarter and R. Amalberti (Eds.), *Cognitive Engineering in the Aviation Domain*, Erlbaum, Hillsdale NJ.
- Cordesman, A. H. and Wagner, A. R. (1996). *The Lessons of Modern War, Vol.4: The Gulf War*, (Boulder, CO: Westview Press).
- Dekker, S. W. A. (1996). Cognitive complexity in management by exception: Deriving early human factors requirements for an envisioned air traffic management world. In D. Harris (Ed.), *Engineering psychology and cognitive ergonomics, Vol. 1*. Aldershot, UK: Ashgate.
- Dekker, S. W. A. and Hollnagel, E., eds. (1999). *Coping with computers in the cockpit*. Aldershot, UK: Ashgate.
- Dekker, S. W. A. and Woods, D. D. (1999). To intervene or not to intervene: The dilemma of management by exception. *Cognition, Technology and Work*, 1, 86-96.
- DeGeus, A. (1988). Planning as Learning. *Harvard Business Review*, March-April, pp. 70-78.
- Flores, F., Graves, M., Hartfield, B. and Winograd, T. (1988). Computer systems and the design of organizational interaction. *ACM Transactions on Office Information Systems*, 6, 153-172.
- Hirschhorn, L. (1997). Quoted in Cook, R. I., Woods, D. D. and Miller, C. (1998). *A Tale of Two Stories: Contrasting Views on Patient Safety*. National Patient Safety Foundation, Chicago IL, April 1998 (available at www.npsf.org).
- Hollnagel, E. and Woods, D. D. (1983). Cognitive systems engineering: New wine in new bottles. *Int. J. Man-machine studies*, 18, 583-600.
- Hoffman, R. and Woods, D. D. (2000). Studying Cognitive Systems in Context. *Human Factors*, 42(1), 1-7.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Klein, G. A., Orasnu, J., Calderwood, R., and Zsombok, C. E. (Eds.) (1993). *Decision making in action: Models and methods*. Norwood, NJ: Ablex.
- Norman, D. A. (1988). *The Psychology of Everyday Things*. New York, NY: Basic Books.
- Obradovich, J. and Woods, D. D. (1996). Users as designers: How people cope with poor HCI design in computer-based medical devices. *Human Factors*, 38(4), 574-592.
- Rasmussen, J. (1986). Information processing and human-machine interaction: An approach to cognitive engineering. Amsterdam: North-Holland.

Rasmussen, J. and Lind M. (1981). *Coping with complexity* (Risø-M-2293). Risø National Laboratory, Roskilde, Denmark: Electronics Department.

Rasmussen, J. (2000). Human Factors in a Dynamic Society. *Ergonomics*, 43, 869—879.

Roth, E.M. and Mumaw, R.J. (1995). Using cognitive task analysis to define human interface requirements for first-of-a-kind systems. Proceedings of the Human Factors and Ergonomics Society 39th Annual Meeting, San Diego, CA, Oct. 9-13, 1995. (pg 520-524).

Sanders, E. B.-N. (2000). Generative Tools for Co-Designing. In *Collaborative Design*, Scrivener, Ball and Woodcock (eds.). London: Springer-Verlag.

Sarter, N. and Amalberti, R., eds. (2000). *Cognitive Engineering in the Aviation Domain*, Erlbaum, Mahwah NJ.

Sarter, N. B., Woods, D. D., and Billings, C. E. (1997). Automation surprises. In G. Salvendy (Ed.) *Handbook of human factors/ergonomics, second edition*. New York, NY: Wiley.

Smith, P., Woods, D., McCoy, E., Billings, C., Sarter, N., Denning R and Dekker, S. (1998). Using forecasts of future incidents to evaluate future ATM system designs. *Air Traffic Control Quarterly*, 6(1), 71—85.

Winograd, T. and Flores, F. (1986). *Understanding computers and cognition*. Norwood, NJ, Ablex.

Weick, K. (1979). *The Social Psychology of Organizing*. New York: Random House.

Woods, D. D. (1988). Coping with complexity: The psychology of human behavior in complex systems. In L.P. Goodstein, H.B. Andersen, and S.E. Olsen, editors, *Mental Models, Tasks and Errors*, Taylor & Francis, London (p. 128-148).

Woods, D. D. (1998). Designs are hypotheses about how artifacts shape cognition and collaboration. *Ergonomics*, 41, 168-173.

Woods, D. D., Johannesen, L. J., Cook, R. I., and Sarter, N. B. (1994). *Behind human error: Cognitive systems, computers and hindsight*. Dayton, OH: CSERIAC.

Woods, D. D., Patterson, E. S., Corban, J. and Watts, J.C. (1996). Bridging the Gap between User-Centered Intentions and Actual Design Practice. Proceedings of the 40th Annual Meeting of the Human Factors and Ergonomics Society, September, 1996.

Woods, D. D. and Tinapple, D. (1999). W³: Watching Human Factors Watch People at Work. Presidential Address, 43rd Annual Meeting of the Human Factors and Ergonomics Society, September 28, 1999. Multimedia Production at <http://cse1.eng.ohio-state.edu/hf99/>

Woods, D. D., Christoffersen, K. and Tinapple, D. (2000). Complementarity and Synchronization as Strategies for Practice-Centered Research and Design. Plenary Address, 44th Annual Meeting of the Human Factors and Ergonomics Society and International Ergonomic Association, August 1, 2000. Multimedia Production available at <http://cse1.eng.ohio-state.edu/iea2000/>